

Bar_{QL}: Collaborating through Change

Oliver Kennedy and Lukasz Ziarek
SUNY Buffalo
{okennedy, lziarek}@buffalo.edu

ABSTRACT

Applications such as Google Docs, Office 365, and Dropbox show a growing trend towards incorporating multi-user *live collaboration* functionality into web applications. These collaborative applications share a need to efficiently express shared state, and a common strategy for doing so is a shared log abstraction. Extensive research efforts on log abstractions by the database, programming languages, and distributed systems communities have identified a variety of optimization techniques based on the algebraic properties of updates (*i.e.*, pairwise commutativity, subsumption, and idempotence). Although these techniques have been applied to specific applications and use-cases, to the best of our knowledge, no attempt has been made to create a general framework for such optimizations in the context of a non-trivial update language. In this paper, we introduce mutation languages, a low-level framework for reasoning about the algebraic properties of state updates, or mutations. We define Bar_{QL}, a general purpose state-update language, and show how mutation languages allow us to reason about the algebraic properties of updates expressed in Bar_{QL}.

1. INTRODUCTION

Over the past several years, many web applications have been released that duplicate and improve on the functionality of desktop applications (*e.g.* Google Docs). A natural consequence of this shift from the desktop to the web is that applications have become more collaborative. Fully featured word processors, presentation editors, spreadsheets, and drawing programs now exist that allow users to collaboratively edit, view, and annotate documents in “real-time.” Although these *collaborative applications* are structured using a client/server model, the core functionality of the application is typically built into the client. The server’s primary role is solely to relay state updates between clients. In spite of this apparent structural simplicity, collaborative application developers must still expend substantial effort to build scalable and efficient infrastructures for their applications. To address this concern, we present the theoretical foundations for a generalized server infrastructure for collabora-

tive applications: Laasie¹. Laasie’s primary goal is to encode and replicate application state through a distributed log datastructure. Clients perform changes to application state by appending them to the log.

The primary motivation for this design is to allow clients to easily recover from link failures (*e.g.*, when the host platform changes networks or after it wakes from sleep mode) by maintaining a pointer to the most recent log entry that they have seen. The server can bring a client up to the most recent state by replaying all log entries that appear after the client’s pointer.

Crucially, updates are expressed in the log in terms of *intent* rather than *effect*. Below, we introduce and discuss Bar_{QL}, an update language that can express conditionals and iteration over complex hierarchical datatypes. Updates expressed in Bar_{QL} are not evaluated, but rather appended as-is to the log. This simplifies the semantics of out-of-order appends and makes it easier to express updates as increments (*i.e.*, deltas) rather than fixed write operations (*e.g.*, `var := 3`).

In short, Bar_{QL} allows the operational semantics of updates to be managed as first class data objects.

Although an append-only log is a useful high-level abstraction, in practice it becomes necessary to compact the log to bound its size. For example, a snapshot of the application state can be substituted for all log entries that precede it. Unfortunately, eliminating all log entries preceding the snapshot also invalidates all clients at states preceding the snapshot as well. These clients must be (effectively) restarted from scratch, negating the benefits of a log.

In this paper, we present a general framework for reasoning about log updates. We consider two properties of each rewrite: (1) *Correctness*, or whether the rewritten log updates collectively generate a state identical to the original sequence of updates, and (2) *Recoverability*, or whether the rewritten log can be used to bring a client at any state up to the most recent state. We then proceed to show how to define *incremental* deletion and composition rewrites of the log and provide realistic “real-world” bounds on their behavior. This is accomplished through the definition and use of mutation languages in the following sections.

The contributions of this paper are as follows:

¹Log-As-A-Service Infrastructure

1. The design and formalism of *mutation languages*, a general framework for reasoning about the correctness and recoverability of log rewrites, and an analysis of the computational complexity of doing so.
2. The construction of mutation languages for composite *hierarchical datatypes* derived from mutation languages for simpler primitive types.
3. The formal definition of a log-based update language named $\text{Bar}_{Q\mathcal{L}}$.
4. A reduction from $\text{Bar}_{Q\mathcal{L}}$ to a composite mutation language, and computational complexity result for computing the correctness and recoverability of log rewrites for $\text{Bar}_{Q\mathcal{L}}$.
5. An incremental algorithm for identifying candidate log rewrites belonging to two rewrites classes: deletion and composition, with amortized constant time complexity.

1.1 Roadmap

Our ultimate goal in this paper is to demonstrate the construction of a practical log rewrite oracle for a non-trivial update language for composite types. For any given rewrite of a log, this oracle will determine both the correctness and recoverability of the rewrite. Before defining the oracle we first define in Section 2 a specification of a nontrivial update language ($\text{Bar}_{Q\mathcal{L}}$). We then use this language to formalize the notion of update logs and log-rewrites, and provide formal definitions of the correctness and recoverability of a log rewrite.

In Section 3 we formally define the mutation and mutation language abstractions. A mutation is simply an expression of *change* and a mutation family is a collection of mutations with properties (e.g., commutativity). We also identify two binary operations (**merge** and **compose**) over mutations in a mutation language that we will use to simplify the translation of $\text{Bar}_{Q\mathcal{L}}$ update queries into equivalent mutations.

Section 3.1 outlines the construction of a log rewrite oracle for any mutation language. This construction is based on language-specific oracles that evaluate algebraic properties of updates (Commutativity, Subsumption, and Idempotence).

In Section 4 we define a mutation language $\bar{\mathcal{L}}$ (LBar), and show a reduction from $\text{Bar}_{Q\mathcal{L}}$ to $\bar{\mathcal{L}}$. We provide definitions of **merge**, **compose**, as well as impractical definitions of the algebraic property oracles. Using a $\bar{\mathcal{L}}$, we define a practical set of algebraic property oracles that allow us to construct a log rewrite oracle for $\text{Bar}_{Q\mathcal{L}}$.

2. HIGH LEVEL SEMANTICS

In this section we introduce $\text{Bar}_{Q\mathcal{L}}$, a log-based update language loosely based on the Monad Algebra [25] with unions and aggregates. Unlike Monad Algebra, which uses sets as

the base collection type, $\text{Bar}_{Q\mathcal{L}}$ uses maps² and has weaker type semantics along the lines of [11]. Furthermore, $\text{Bar}_{Q\mathcal{L}}$ is intentionally limited to operations with linear computational complexity in the size of the input data; neither the pairwith nor cross-product operations of Monad Algebra are included. In our domain, this is not a limitation, as the server is acting primarily as a relay for state. Full cross-products can be transmitted to clients more efficiently in their factorized form, and each client is expected to be capable of computing cross products locally³.

The domains and grammar for $\text{Bar}_{Q\mathcal{L}}$ are given in Fig. 1. We use C to range over constants, p over primitives (strings, integers, floats, and booleans), k over keys, Q over queries, τ over types, v over values of type τ , and θ over binary operations over primitive types. The type τ operated over by $\text{Bar}_{Q\mathcal{L}}$ queries is identical to the labeled trees of [11], and is equivalent to unstructured XML or JSON. Values are either of primitive type, null, or collections (mappings from k to τ). Note that collections are total mappings; for instances, a singleton can be defined as the collection where all keys except one map to the *null* value. By convention, when referring to collections we will implicitly assume the presence of this mapping for all keys that are not explicitly specified in the rules themselves.

We formalize $\text{Bar}_{Q\mathcal{L}}$ in Fig. 2 in terms of a big-step operational semantics. Order of evaluation is defined by the structure of the rules.

In $\text{Bar}_{Q\mathcal{L}}$ queries are monads, structures that represent computation. Reducing the query corresponds to evaluating the computation expressed by that query. The rules for *PrimitiveConstant*, *Null* and *EmptySet* all defined operations take an input value and produce a constant value regardless of input. The rule *PrimitiveConstant* produces a primitive constant c , the rule *Null* produces the *null* value, and the rule *EmptySet* produces an empty set. We define an empty set as a collection that is a total mapping where all keys map to the *null* value, represented as: $\{ * \rightarrow \text{null} \}$. The *Identity* operation passes through the input value unchanged. *Subscripting* and *Singleton* are standard operations. In comparison to Monad Algebra, these operations correspond to not only the singleton operation over sets, but also the tuple constructor and projection operations. Because collection elements are identified by keys, we can reference specific elements of the collection in much the same way as selection from a tuple. The most significant way in which $\text{Bar}_{Q\mathcal{L}}$ differs from Monad Algebra is its use of the *Merge* operation (\Leftarrow) instead of set union (\cup). \Leftarrow combines two sets, overwriting undefined entries (keys for which the collection maps to *null*) with their values from the other collection.

$$(\{A := 1\} \Leftarrow \{B := 2\})(\text{null}) = \{A \rightarrow 1, B \rightarrow 2\}$$

²Maps are also popularly referred to as hashes, dictionaries, or lookup tables.

³Joins are an area of concern however, and future work will consider extensions to $\text{Bar}_{Q\mathcal{L}}$ for this purpose.

$c \in \text{Constant} : \rightarrow p$	$k \in \text{Key}$	$Q := Q.k \mid Q \Leftarrow Q \mid \text{map } Q \text{ using } Q \mid Q \text{ op}_{[\theta]} Q$
$p \in \text{Primitive}$	$Q \in \text{Query} : \tau \rightarrow \tau$	$\mid \text{agg}_{[\theta]}(Q) \mid \text{agg}_{[\Leftarrow]}(Q) \mid \text{filter } Q \text{ using } Q$
$v \in \text{Value}$	$\tau \in \text{Type} : p \mid \{k_i \rightarrow \tau_i\} \mid \text{null}$	$\mid \text{if } Q \text{ then } Q \text{ else } Q \mid Q \circ Q \mid \mathbf{c} \mid \text{null} \mid \emptyset$
$\theta \in \text{BinaryOp}$		

Figure 1: Domains and grammar for $\text{Bar}_{Q\mathcal{L}}$.

If a key is defined in both collections, the right collection takes precedence.

$$(\{A := 1\} \Leftarrow \{A := 2\})(\text{null}) = \{A \rightarrow 2\}$$

The merge operator can be combined with singleton and identity to define updates to collections:

$$(\mathbf{id} \Leftarrow \{A := 3\})(\{A \rightarrow 1, B \rightarrow 2\}) = \{A \rightarrow 3, B \rightarrow 2\}$$

Subscripting can be combined with merge, singleton, and identity to define point modifications to collections.

$$(\mathbf{id} \Leftarrow \{A := (\mathbf{id}.A \Leftarrow \{B := 2\})\})(\{A \rightarrow \{C \rightarrow 1\}\}) \\ = \{A \rightarrow \{B \rightarrow 2, C \rightarrow 1\}\}$$

Primitive binary operators are defined monadically with operation *PrimBinOp*, and include basic arithmetic, comparisons, and boolean operations. These operations can be combined with identity, singleton, and merge to define updates. For example, to increment *A* by 1, we write

$$\{\mathbf{id} \Leftarrow \{A := \mathbf{id}.A + 1\}\}(\{A \rightarrow 2\}) = \{A \rightarrow 3\}$$

$\text{Bar}_{Q\mathcal{L}}$ provides constructs for mapping, flattening and aggregation. The *Map* operation is analogous to its definition in Monad Algebra, save that key names are preserved. The *Flatten* operation is also similar, except that it uses \Leftarrow , instead of \cup as in Monad Algebra. The *PrimitiveAggregation* class of operators defines aggregation using any closed binary operator θ operating over primitive type.

To increment all children of the root by 1 we write:

$$(\text{map id using } (\mathbf{id} + 1))(\{A \rightarrow 1, B \rightarrow 2\}) = \{A \rightarrow 2, B \rightarrow 3\}$$

To increment the child *C* of each child of the root by 1, we write

$$(\text{map id using } (\mathbf{id} \Leftarrow \{C := \mathbf{id}.C + 1\}))(\{A \rightarrow \{C \rightarrow 1\}, B \rightarrow \{C \rightarrow 2, D \rightarrow 1\}\}) \\ = \{A \rightarrow \{C \rightarrow 2\}, B \rightarrow \{C \rightarrow 3, D \rightarrow 1\}\}$$

Finally, $\text{Bar}_{Q\mathcal{L}}$ supports *Conditionals* and *Filtering*, as well as *Composition* of queries.

3. MUTATION LANGUAGES

We will now temporarily step back from $\text{Bar}_{Q\mathcal{L}}$ in order to refine our understanding of *update logs*. At its simplest, an update log encodes a state value as a sequence of state mutations applied iteratively, first to a default “empty” state, and then to the output of the prior transformation.

If the fundamental primitive of an update log is the state transformation, then the fundamental operation is composition of state mutations. As a basis for reasoning about the safety properties of changes to this log, we begin with an outline for simple algebras over the composition of state mutations.

DEFINITION 1. A *mutation* is an arbitrary transformation $M : \tau \mapsto \tau$ mapping values of some state type τ to new values of the same type. A mutation may be parameterized by an set of additional values R . We write such a mutation as $M_R(v)$. We say the mutation $M_R(v)$ is:

- ... **destructive** if M_R is independent of v .
- ... **idempotent** if $\forall v, R : M_R(v) \equiv M_R(M_R(v))$

EXAMPLE 1. Consider an application that encodes its state as a single integer (i.e., $\tau = \mathbb{Z}$). Such an application might employ the two mutations “replace by 0”, and “increment by 1”:

$$M_{:=0}(x) \mapsto 0 \quad M_{++}(x) \mapsto x + 1$$

The replace operation is both destructive and idempotent. The increment operation is neither.

We can use parameters to create families of mutations. For example, we can use a single parameter Y to define a family of mutations “replace by Y ” ($M_{:=Y}$), or “increment by Y ” ($M_{+=Y}$).

Having defined mutations in the abstract as functions, we can now formally define the abstract composition of mutations as simple left-first function composition.

$$(M \circ M')(x) \equiv M'(M(x))$$

PROPOSITION 1. Composition is associative.

PROOF. By Equivalence

$$((M \circ M') \circ M'')(x) \equiv M''(M'(M(x))) \equiv (M \circ (M' \circ M''))(x)$$

□

We can define a composition algebra for any set of mutations \vec{M} with identical kinds. We consider two properties in this algebra: (1) pairwise commutativity and (2) subsumption. Unlike the traditional algebraic notion of commutativity, we consider only the pairwise commutativity of individual mutations. That is, instead of saying that \circ is commutative, we say that M and M' *commute* iff $(M \circ M') \equiv (M' \circ M)$. Subsumption is also defined pairwise; we say that M' *subsumes* M iff $M \circ M' \equiv M'$.

$$\begin{array}{c}
\text{PrimitiveConstant} \quad \frac{}{\widehat{c}(v) \mapsto c} \quad \text{Null} \quad \frac{}{\widehat{null}(v) \mapsto null} \quad \text{EmptySet} \quad \frac{}{\widehat{\emptyset}(v) \mapsto \{ * \rightarrow null \}} \quad \text{Identity} \quad \frac{}{\text{id}(v) \mapsto v} \\
\\
\text{Subscripting} \quad \frac{Q(v) \mapsto \{ \dots, k \rightarrow r, \dots \}}{(Q.k)(v) \mapsto r} \quad \text{Singleton} \quad \frac{Q(v) \mapsto r}{\{key := Q\}(v) \mapsto \{k \rightarrow r, * \rightarrow null\}} \\
\\
\text{Merge} \quad \frac{Q_1(v) \mapsto \{k_i \rightarrow r_i\} \quad Q_2(v) \mapsto \{k_j \rightarrow r_j\}}{(Q_1 \widehat{\Leftarrow} Q_2)(v) \mapsto \{k \rightarrow r \mid (k = k_i = k_j) \wedge (((r = r_i) \wedge (r_j = null)) \vee ((r = r_j) \wedge (r_j \neq null)))\}} \\
\\
\text{Map} \quad \frac{Q_{coll}(v) \mapsto \{k_i \rightarrow v_i\} \quad Q_{map}(v_i) \mapsto r_i}{(\text{map } Q_{coll} \text{ using } Q_{map})(v) \mapsto \{k_i \rightarrow r_i \mid v_i \neq null\}} \quad \text{PrimBinOp} \quad \frac{Q_1(v) \mapsto r_1 : p \quad Q_2(v) \mapsto r_2 : p \quad \theta \in \{+, *, -, /, =, \text{AND}, \text{OR}, \neq, <, \leq, >, \geq\}}{(Q_1 \widehat{\theta} Q_2)(v) \mapsto r_1 \theta r_2} \\
\\
\text{Flatten} \quad \frac{Q_{coll}(v) \mapsto \{k_i \rightarrow v_i\}}{(\text{agg}_{[\Leftarrow]}(Q_{coll}))(v) \mapsto (v_0 \Leftarrow v_1 \Leftarrow \dots)} \quad \text{PrimitiveAggregate} \quad \frac{Q_{coll}(v) \mapsto \{k_i \rightarrow v_i\}}{(\text{agg}_{[0]}(Q_{coll}))(v) \mapsto (((v_0 \theta v_1) \theta v_2) \theta \dots)} \\
\\
\text{IfThenElse} \quad \frac{Q_{cond}(v) \mapsto true \quad Q_{then}(v) \mapsto r_{then}}{(\text{if } Q_{cond} \text{ then } Q_{then} \text{ else } Q_{else})(v) \mapsto r_{then}} \quad \frac{Q_{cond}(v) \mapsto false \quad Q_{else}(v) \mapsto r_{else}}{(\text{if } Q_{cond} \text{ then } Q_{then} \text{ else } Q_{else})(v) \mapsto r_{else}} \\
\\
\text{Filter} \quad \frac{Q_{coll}(v) \mapsto \{k_i \rightarrow v_i\} \quad Q_{cond}(v_i) \mapsto t_i}{(\text{filter } Q_{coll} \text{ using } Q_{cond})(v) \mapsto \{k_i \rightarrow v_i \mid t_i \wedge v_i \neq null\}} \quad \text{Composition} \quad \frac{Q_1(v) \mapsto r_1 \quad Q_2(r_1) \mapsto r_2}{(Q_1 \circ Q_2)(v) \mapsto r_2}
\end{array}$$

Figure 2: A formal operational semantics for $\text{Bar}_{Q_{\mathcal{L}}}$.

DEFINITION 2. A mutation language \mathcal{L} is the 4-tuple: $\langle \tau, \vec{M}, \mathcal{S}, \mathcal{C} \rangle$ consisting of:

1. A state type τ
2. A set of mutations \vec{M} of kind $\tau \mapsto \tau$. This set must include the identity mutation $\text{id}(x) \mapsto x$.
3. A binary relation $\mathcal{S}(M, M')$ that holds if M is subsumed by M' .
4. A symmetric binary relation $\mathcal{C}(M, M')$ that holds if M commutes with M' .

We will use the shorthand $\mathcal{S}(M) \equiv \mathcal{S}(M, M)$ to denote the unary idempotence relation.

A mutation language encapsulates the composition algebra for a specific set of mutations, together with a set of rules for determining idempotence, pairwise commutativity, and subsumption on mutations in the language.

EXAMPLE 2. On simple mutation languages, these properties can be determined quite efficiently. For the mutation language defined from the mutation language families in Example 1 ($M_{:=Y}$ and $M_{+=Y}$), we can define the commutativity and subsumption relations by simple structural tests on the mutations being related: $\mathcal{C}(M_{:=Y}, M_{:=Y})$, $\mathcal{C}(M_{+=Y}, M_{+=Y'})$, and $\mathcal{S}(M, M_{:=Y})$ are the only relations that hold. The identity mutation for this language is $M_{+=0}$.

For more complex classes of mutations, this definition can be too strong. Consequently, for the remainder of the paper, we will limit ourselves to *weak mutation languages*, where the relations \mathcal{S}, \mathcal{C} are conservative approximations. If the relation holds then the corresponding property is guaranteed to hold, but not visa versa.

Finally, we will define two notions of closure for a mutation language: First, a mutation language \mathcal{L} is closed over composition if the composition of two mutations $M, M' \in \mathcal{L}$ is also in \mathcal{L} .

$$\forall M, M' \in \mathcal{L} : \exists M'' \equiv (M \circ M') \in \mathcal{L}$$

Second, a mutation language \mathcal{L} is closed over binary operation $\theta : \tau \times \tau \mapsto \tau$ if there exists a mutation in \mathcal{L} that computes the result of applying θ to the output of mutations $M, M' \in \mathcal{L}$.

$$\forall M, M' \in \mathcal{L} : \exists M'' \in \mathcal{L} : M''(x) \equiv (M(x) \theta M'(x))$$

EXAMPLE 3. Our toy mutation language from Example 2 can be shown to be closed over composition, addition, and subtraction, but not multiplication. The details of this proof are left to the reader.

We will use the two binary operations **compose** and **merge_θ** to denote the result of combining two mutations by composition or by binary operation θ (respectively), for any mutation language closed over composition or θ (respectively). Note

that the existence of either function provably demonstrates the corresponding type of closure.

3.1 Mutation Logs

We now turn to our primary subject: logs. Our goal in this section is to develop formalisms, first for the logs themselves, and second for reasoning about how the logs can be transformed, or rewritten, while preserving certain critical properties.

A *log* is a sequence of updates to an application's state, expressed as a numbered sequence of mutations: M_1, \dots, M_n . A log defines a corresponding sequence of application states: v_0, \dots, v_n . We obtain state v_i by starting with a default state v_0 , and applying mutations M_1, \dots, M_i in order. In other words, for a mutation language closed under composition, v_i is the result of composing the first i mutations in the log.

$$v_i = (M_1 \circ \dots \circ M_i)(v_0)$$

We refer to the subscript of a state or mutation as its *timestamp* (i.e., v_i and M_i have timestamp x). We define the *current state* of a log of size n to be the state v_n . The current state can be *recovered* from any intermediate state v_i by applying the composition of all mutations after x .

$$v_n = (M_{x+1} \circ \dots \circ M_n)(v_i)$$

Recovery is central to the design of Laasie. A client can recover from a transient disconnection by replaying only those mutations that occurred while the client was disconnected, rather than forcing it to reload the full application state from scratch.

3.1.1 Log Rewrites

A log rewrite \mathcal{R} is defined generally as an operation that transforms one sequence of mutations M_1, \dots, M_n into a new sequence M'_1, \dots, M'_n .

Because of our interest in recovery, we are interested in preserving a correspondence between timestamps in the pre- and post-rewritten states v_i and v'_i (respectively). Consequently, we will assume that each pre-rewrite state corresponds to the post-rewrite state with the same timestamp. Note that this limits us to size-preserving rewrites. As we will soon see, this can be done without loss of generality. We specifically consider two classes of size-preserving log rewrites: *delete* and *compose*.

Delete. We can effect a size-preserving deletion rewrite by replacing the deleted mutation with the no-op identity operation (**id**). The rewrite $\mathcal{R}_{\text{del}}(x)$, which deletes mutation M_x is defined as

$$M'_i = \begin{cases} M_i & \dots & i \neq x \\ \text{id} & \dots & i = x \end{cases}$$

Compose. For a mutation language closed over composition, we can merge two mutations into the log into a single log entry. The log size is preserved by inserting an **id** mutation. For reasons that will soon become clear, the composite

mutation replaces the mutation with the higher timestamp, and the inserted **id** replaces the mutation with the lower timestamp. The rewrite $\mathcal{R}_{\text{cmp}}(x, y)$, which merges mutations M_x and M_y is defined as

$$M'_i = \begin{cases} M_i & \dots & i \notin \{x, y\} \\ \text{id} & \dots & i = x \\ M_x \circ M_y & \dots & i = y \end{cases}$$

3.1.2 Rewrite Properties

Now that we have defined log rewrites, we begin to consider what constitutes a legitimate log rewrite. We define three correctness properties for log rewrites: *tail-correctness*, *recoverability*, and *\bar{r} -recoverability*. We will also show how to use the subsumption and commutativity relations of a mutation language \mathcal{S} , \mathcal{C} to determine when these properties are guaranteed to be satisfied, independent of data, for a delete, compose, or commute rewrite.

Tail-Correctness. We start with the simplest of the log-rewrite properties.

DEFINITION 3. A log rewrite is *tail-correct* if the current state v_n of the log is identical to the current state v'_n of the rewritten log. That is:

$$(M_1 \circ \dots \circ M_n)(v_0) = (M'_1 \circ \dots \circ M'_n)(v_0)$$

LEMMA 1. The rewrite $\mathcal{R}_{\text{del}}(x)$ is *tail-correct* if M_x is subsumed by the aggregate composition of all mutations following it: $\mathcal{S}(M_x, (M_{x+1} \circ M_{x+2} \circ \dots \circ M_n))$.

PROOF. The identity operation has no effect on the state, and can be inserted anywhere. By subsumption, we have that

$$M_x \circ \dots \circ M_n \equiv M_{x+1} \circ \dots \circ M_n$$

Thus, $v_n = v'_n$ □

LEMMA 2. The rewrite $\mathcal{R}_{\text{cmp}}(x, y)$ is *tail-correct* for any mutation language closed over composition if M_x commutes with the aggregate composition of all mutations between it and M_y : $\mathcal{C}(M_x, (M_{x+1} \circ \dots \circ M_{y-1}))$

PROOF. As before, identity has no effect on the state. If $x = y - 1$, then the merged mutations is equivalent to the separate mutations by Proposition 1. Otherwise, by commutativity, we have that

$$M_x \circ \dots \circ M_{y-1} \equiv M_{x+1} \circ \dots \circ M_{y-1} \circ M_x$$

Once M_x and M_y are adjacent, they can be merged just as before. □

EXAMPLE 4. Consider our toy mutation language from Example 2. From the subsumption relation \mathcal{S} , we can infer that it is *tail-correct* to delete any mutation preceding a replace mutation ($M_{:=y}$).

From the commutativity relation \mathcal{C} , we can infer that it is *tail-correct* to merge any two mutations in an unbroken sequence of increment mutations ($M_{+=y}$), or to merge a replace mutation with its immediate successor.

Recoverability. Although tail-correctness provides a useful baseline for further discussions of log rewrites, it only takes a single state: the current state into consideration. As such, it fails to capture any of the benefits of having a log in the first place. We now consider a property that is strictly stronger than tail-correctness, and which allows us to reason about the possibility of recovery from any intermediate state. We start with a per-timestamp notion of recoverability

DEFINITION 4. A log rewrite is recoverable from timestamp i (or equivalently state v_i) if the final state v_n of the original log can be obtained by applying the sequence of rewritten mutations following timestamp i to the state v_i , taken from the original log.

$$(M_1 \circ \dots \circ M_n)(v_0) = (M'_{i+1} \circ \dots \circ M'_n)(v_i)$$

Or equivalently (because v_i is defined by the original log)

$$(M_1 \circ \dots \circ M_n)(v_0) = (M_1 \circ \dots \circ M_i \circ M'_{i+1} \circ \dots \circ M'_n)(v_0)$$

DEFINITION 5. A log rewrite is recoverable if it is recoverable from all timestamps in the log (i.e., $i \in [0, n]$)

Note that tail-correctness is the special case of recoverability from timestamp 0.

LEMMA 3. If the log rewrite $\mathcal{R}_{\text{del}}(x)$ is tail-correct, it is recoverable

PROOF. Recoverability from any state v_i s.t. $i < x$ is equivalent to tail-correctness, because these states are unaffected by the rewrite. Recoverability when $i \geq x$ is guaranteed always: The state v_i being recovered from is taken before the rewrite, and mutations M'_{x+1}, \dots, M'_n are identical to their pre-rewrite counterparts. \square

This proof shows a tight coupling between correctness and recoverability, and illustrates an intriguing log partitioning. If a rewrite only modifies mutations that fall within a fixed range, recoverability “errors” can only occur at states that fall within that same range.

PROPOSITION 2. Let \mathcal{R} be a tail-correct log rewrite, which only alters log entries at timestamps in the range $[x, y]$. Mutations outside of this range are unaffected by \mathcal{R} . \mathcal{R} is recoverable iff it is recoverable from all states $v_i \in [x, y]$

PROOF. The proof is identical to that of Lemma 3. \square

LEMMA 4. The rewrite $\mathcal{R}_{\text{cmp}}(x, y)$ is recoverable if it is correct, and if M_x is idempotent: $\mathcal{S}(M_x, M_x)$

PROOF. From the commutativity property required to show correctness, we have that $M_x \circ \dots \circ M_{y-1} \equiv M_{x+1} \circ \dots \circ M_{y-1} \circ M_x$. For all $i \geq x$, state $v_i = (M_x \circ \dots \circ M_i)(v_{x-1})$. Thus, $(M'_{i+1} \circ M'_y)(v_i) \equiv (M_x \circ \dots \circ M_x \circ M_y)(v_x)$. By commutativity, we can rewrite this expression as $M_{x+1} \circ \dots \circ M_x \circ M_x \circ M_y$. By idempotence, this is equivalent to the original rewritten expression, and by Proposition 2 the proof devolves to that of correctness.

EXAMPLE 5. Returning to the toy mutation language from Example 2, we see that although it is tail-correct to merge any two increment mutations, it is not recoverable.

Consider the log $(M_{:=1}, M_{+=2}, M_{+=3})$. After applying the rewrite $\mathcal{R}_{\text{cmp}}(2, 3)$, we get $(M_{:=1}, \text{id}, M_{+=5})$. After the rewrite, it is no longer possible to recover from state $v_2 (= 3)$, as the mutation $M_{+=2}$ would effectively be applied twice.

\vec{t} -recoverability. The intent of recoverability is to protect disconnected clients from reaching an inconsistent state when log entries are replayed. However, to guarantee full recoverability, we must discard many potentially useful log rewrites. In a practical setting, a server will not need to guarantee recoverability for all timestamps.

DEFINITION 6. Given a set of timestamps \vec{t} , a log rewrite is \vec{t} -recoverable if it is recoverable from every $t \in \vec{t}$.

By tracking when clients disconnect (regardless of whether or not the disconnection is transient), the server can identify ranges of log entries over which non-recoverable log rewrite can still be performed.

THEOREM 1. Let \mathcal{R} be a tail-correct, but non-recoverable log rewrite. Let $[x, y]$ be the minimal range of timestamps affected by \mathcal{R} . \mathcal{R} is \vec{t} -recoverable iff $(\vec{t} \cap [x, y]) = \emptyset$.

PROOF. Follows from Proposition 2

4. REDUCING $\text{Bar}_{Q\mathcal{L}}$ TO LBAR

We now apply the principles of mutation languages to $\text{Bar}_{Q\mathcal{L}}$ by constructing a weak mutation language $\tilde{\mathcal{L}}$ (LBar) built around $\text{Bar}_{Q\mathcal{L}}$. Roughly speaking, this mutation language allows a single monolithic $\text{Bar}_{Q\mathcal{L}}$ query to be subdivided into a set of disjoint operations, each applied to a specific point in the path hierarchy. This allows us to easily identify the write dependencies of a $\text{Bar}_{Q\mathcal{L}}$ query at their finest granularity.

We then transform each subdivided operation into a delta form, with a $\text{Bar}_{Q\mathcal{L}}$ query that computes a delta value and a merge operator, a binary function that defines how the delta value is to be merged with the prior state. This update operator simplifies the task of determining commutativity and subsumption at a fine granularity.

We also identify the set of points in the path hierarchy that each query reads from. This set of points forms the set of read dependencies of the query.

Finally, we use the sets of write dependencies, read dependencies, and update operators to efficiently compute the commutativity and subsumption relations \mathcal{C}, \mathcal{S} for a $\text{Bar}_{Q\mathcal{L}}$ query.

4.1 LBar

The typesystem of $\tilde{\mathcal{L}}$ is identical to that of $\text{Bar}_{Q\mathcal{L}}$. To recap: values can be of any primitive type, or a collection, which is a mapping from key names of abstract type k to values. Collections can be organized into a hierarchy. We use ϕ to

denote an ordered sequence of key names that defines a *path* through the collection hierarchy.

Point mutations form the basis of $\bar{\mathcal{L}}$, and express updates to individual paths in a $\text{Bar}_{Q\mathcal{L}}$ hierarchy. A *point mutation* is a 3-tuple $\langle \phi, Q, (\theta \mid \emptyset) \rangle$, where ϕ is the path being updated and Q is a $\text{Bar}_{Q\mathcal{L}}$ expression that computes an *update delta* based on the prior state. Every point mutation is annotated with either a binary operation θ , or the *overwrite* annotation \emptyset . The annotation indicates the *combinator* used to merge delta value with the original.

We say that two point mutations are path-disjoint if neither point mutation's path is a prefix of the other's. A *full mutation* in $\bar{\mathcal{L}}$ is a set of pairwise path-disjoint point mutations, which it applies to the state in parallel; The prior state for all point mutations in the set is defined uniformly to be the prior state for the full mutation. Thus, all point mutations are guaranteed to be isolated in the traditional database sense.

As a shorthand, we will use $\omega(\mathcal{M})$ to denote the write set of a full mutation \mathcal{M} , the set of all paths of point mutations in the full mutations:

$$\omega(\mathcal{M}) = \{\phi \mid \langle \phi, Q, \theta \rangle \in \mathcal{M}\}$$

We will also use the shorthand $\mathcal{M}[\phi]$ to denote the point mutation applied to path ϕ for all $\phi \in \omega(\mathcal{M})$.

4.2 Reduction Algorithm

We now present an iterative process for transforming $\text{Bar}_{Q\mathcal{L}}$ expressions into $\bar{\mathcal{L}}$ form. This process begins by creating a full mutation consisting of a single point mutation $\{\langle \emptyset, Q, \emptyset \rangle\}$. The algorithm repeatedly selects an arbitrary point mutation in the set and tries (1) to subdivide point-mutations in this set into finer-grained mutations, and (2) to replace overwrite annotations by extracting binary operations from the point-mutation's query. This process proceeds up to a fixed point.

Operator-Extraction. In their simplest incarnations, both transformations are applied to point queries of the same general form:

$$\langle \phi, (\text{id}.\phi \theta Q'), \emptyset \rangle$$

For a θ that is commutative and associative, any query with a $\text{id}.\phi$ term can be commuted to the front.

In this expression, Q' effectively expresses the delta of the point update, while θ combines it with the original value $\text{id}.\phi$. Consequently, θ becomes the new combinator, and Q' becomes the new update delta.

Key-Extraction. The merge operator (\Leftarrow) is associative (but not commutative). As with binary operators on primitive type, we can compute an update delta of expressions that derive from id . We start by identifying the *change set* of the original query. We start from a point update of the form:

$$\langle \phi, Q, \emptyset \rangle$$

If a query Q returns a value of collection type, its change set $\delta(Q)$ is computed as follows:

- $\delta(\emptyset) = \emptyset$
- $\delta(\{k := Q'\}) = \{k\}$
- $\delta(\text{id}.\phi) = \{*\}$
- $\delta(\text{id}.\phi') = \text{This point mutation can not be subdivided.}$
- $\delta(\text{map } Q' \text{ using } \dots) = \delta(Q')$
- $\delta(\text{filter } Q' \text{ using } \dots) = \delta(Q')$
- $\delta(Q' \Leftarrow Q'') = \delta(Q') \cup \delta(Q'')$
- $\delta(\text{if } \dots \text{ then } Q' \text{ else } Q'') = \delta(Q') \cup \delta(Q'')$
- $\delta(Q' \circ Q'') = \delta Q''[\text{id}/Q']$

The key $*$ is a special key that refers to all keys in the input query input. This special key is treated as a distinct key in the changeset computation. If it is in the changeset for a delta query ($* \in \delta(Q)$), the point mutation modifies the original value (instead of overwriting it), and can be subdivided further as follows.

We begin by generating a delta computation $\Delta_k(Q)$ for each subkey k in the changeset. This includes a delta computation for the special key $*$, which will be applied to all keys in the input that are not explicitly present in the changeset.

- $\Delta_k(\emptyset) = \text{null}$
- $\Delta_k(\{k := Q'\}) = Q'$
- $\Delta_k(\{k' := Q'\}) = \text{null}$
- $\Delta_k(\text{id}.\phi) = \text{id}.\phi.k$
- $\Delta_k(\text{map } Q' \text{ using } Q'') = \Delta_k(Q') \circ Q''$
- $\Delta_k(\text{filter } Q' \text{ using } Q'') =$
 $\text{if } \Delta_k(Q') \circ Q'' \text{ then } \Delta_k(Q') \text{ else null}$
- $\Delta_k(Q' \Leftarrow Q'') =$
 $\text{if } \Delta_k(Q'') \neq \text{null} \text{ then } \Delta_k(Q'') \text{ else } \Delta_k(Q')$
- $\Delta_k(\text{if } \Delta_k(Q) \text{ then } Q' \text{ else } Q'') =$
 $\text{if } \Delta_k(Q) \text{ then } \Delta_k(Q') \text{ else } \Delta_k(Q'')$
- $\Delta_k(Q' \circ Q'') = \Delta_k Q''[\text{id}/Q']$

The resulting expression can be simplified by partial evaluation. In many cases, it will be possible to eliminate operations over *null* values. The result is a set of point mutations, one for each key k in the changeset, including the special key $*$. Once again, $*$ applies to all children at ϕ except those explicitly defined (by being present in the changeset). The resulting set of point mutations is thus defined as

$$\{\langle \phi.k, \Delta_k(Q), \emptyset \rangle \mid k \in \delta(Q) \wedge (\Delta_k(Q) \neq \text{id}.\phi.k)\}$$

Note that we explicitly exclude the identity mutation, as this is effectively a no-op.

4.3 Read Dependencies

We compute the read dependencies of a $\text{Bar}_{Q\mathcal{L}}$ query by first defining a read-normal form for $\text{Bar}_{Q\mathcal{L}}$. We call a query of the form $\text{id}.k_1.k_2.\dots.k_n$ a *point read* at path $\phi = k_1.k_2.\dots.k_n$. A query is in read-normal form if the subscript operator appears only in point reads, or is applied to the special key tmp , defined below. As we now show, any valid query can be transformed into read-normal form:

- $(Q \Leftarrow Q').k \mapsto \text{if } Q'.k \neq \text{null} \text{ then } Q'.k \text{ else } Q.k$
- $(\text{map } Q \text{ using } Q').k \mapsto Q.k \circ Q'$
- $(\text{agg}_{[\Leftarrow]}(Q)).k \mapsto (\text{agg}_{[\Leftarrow]}(\text{map } Q \text{ using } \{ \text{tmp} := \text{id}.k \})).\text{tmp}$
- $(\text{filter } Q \text{ using } Q').k \mapsto \text{if } Q.k \circ Q' \text{ then } Q.k \text{ else null}$
- $(\text{if } Q \text{ then } Q' \text{ else } Q'').k \mapsto \text{if } Q \text{ then } Q'.k \text{ else } Q''.k$
- $(Q \circ Q').k \mapsto Q \circ (Q'.k)$
- $\emptyset.k \mapsto \text{null}$

Given a query Q in read-normal form, we can compute the readset of the query $\rho(Q)$ as follows:

- $\rho(\text{id}.\phi) = \{\phi\}$
- $\rho(Q \Leftarrow Q') = \rho(Q) \cup \rho(Q')$
- $\rho(\text{map } Q \text{ using } Q') = \rho(Q)^4$
- $\rho(Q \circ Q') = \rho(Q) \cup \rho(Q')$
- $\rho(\text{agg}_{[\Leftarrow]}(Q)) = \rho(Q)$
- $\rho(\text{filter } Q \text{ using } Q') = \rho(Q)$
- $\rho(\text{if } Q \text{ then } Q' \text{ else } Q'') = \rho(Q) \cup \rho(Q') \cup \rho(Q'')$
- $\rho(Q \circ Q') = \rho(Q)[\text{id}/Q]$
- $\rho(c|\text{null}|\emptyset) = \emptyset$

4.4 Subsumption and Commutativity

We are now ready to complete the definition of the mutation language 4-tuple for $\bar{\mathcal{L}}$ by defining a conservative approximation of the subsumption and commutativity relations.

Subsumption. A path ϕ is subsumed by a full mutation \mathcal{M} if it or one of its ancestors is *overwritten* by \mathcal{M} , and neither ϕ , nor any of its ancestors or descendants appear in the read set of \mathcal{M} . Abusing syntax, we write this as:

$$\begin{aligned} \mathcal{S}(\phi, \mathcal{M}) \equiv & \\ & (\exists Q, \phi' \in \omega(\mathcal{M}) : (\phi' \sqsubseteq \phi) \wedge (\mathcal{M}[\phi'] = \langle \phi', Q, \emptyset \rangle)) \\ & \wedge (\nexists \phi' \in \rho(\mathcal{M}) : (\phi' \sqsubseteq \phi) \vee (\phi \sqsubseteq \phi')) \end{aligned}$$

Here, \sqsubseteq denotes the ancestor of relation.

⁴This is a conservative approximation.

A mutation \mathcal{M} is subsumed by \mathcal{M}' if all paths in the write set of \mathcal{M} are subsumed by \mathcal{M}' :

$$\mathcal{S}(\mathcal{M}, \mathcal{M}') \equiv \forall \phi \in \omega(\mathcal{M}') : \mathcal{S}(\phi, \mathcal{M}')$$

Commutativity. Two point mutations applied to the same path ϕ , $\langle \phi, Q, \theta \rangle$ and $\langle \phi, Q', \theta' \rangle$ commute iff θ commutes with θ' . Two point mutations applied to different paths, $\langle \phi, Q, \theta \rangle$ and $\langle \phi', Q', \theta' \rangle$ commute iff each of the following conditions holds: (1) ϕ is neither an ancestor, nor descendant of ϕ' , (2) ϕ is neither an ancestor, nor descendant of a path in the read set $\rho(Q')$, and (3) ϕ' is neither an ancestor, nor descendant of a path in the read set $\rho(Q)$.

Two full mutations commute, if all pairs of point mutations commute. Again, abusing syntax:

$$\mathcal{C}(\mathcal{M}, \mathcal{M}') \equiv \forall m \in \mathcal{M}, m' \in \mathcal{M}' : \mathcal{C}(m, m')$$

5. RELATED WORK

There has been much work focused on the formalization of query languages and database models [3, 4, 26]. Much of this work is based on monad algebra, Lawvere theories, and universal algebra [23, 7, 5, 22]. Manes *et al.* [27] showed how to implement collection classes using monads. Cluet [17] is an algebra based query language for an object-oriented database system. Our work is based on the same fundamental theories. In the following we compare our work to previous results.

Languages for Transforming Hierarchical Data. There has been considerable work [12, 2, 3, 1] on the transformation of hierarchical data. Two approaches have become dominant in this area: Nested Relational Calculus [31] and the Monad Algebra [25]. Our own approach is closely based on the latter, adapted for use with labeled sets, and with the intentional exclusion of the superlinear time complexity pair-with operator (or equivalently, the cartesian cross-product).

Semistructured Data. Also closely related is work on managing semistructured data [11]. The vast majority of recent efforts in this area have been on querying and transforming XML data. One formalization by Koch [24] is also closely based on Monad Algebra. Work by Cheney follows a similar vein, in particular (F)LUX [15, 16], a functional language for XML updates. In [8], Benedikt and Cheney present a formalism for synthesizing the output schema of XML transformations, similar to our notion of the compositional compatibility of mutations. More recently, there has also been interest in querying lighter-weight semistructured data representations like JSON[9, 10].

Algebraic Properties of State Updates. The distributed systems community has identified a number of algebraic properties of state mutations that are useful in distributed concurrency control. Commutativity of updates has been explored extensively [34, 32], but the typical assumption is that a domain-specific commutativity oracle is available, such as for edits to textual data [32, 28]. Our notion of subsumption

is quite similar to the Badrinath and Ramamritham [6]’s recoverability property. Unlike subsumption, this property is defined in terms of observable side-effects rather than state, but is otherwise identical. Like prior work on commutativity, they assume that a domain-specific oracle has been provided. Several efforts have been made to understand domain-specific reconciliation strategies. Feldman *et al.*’s Operational Transforms [21] are analogous to our mutation languages, but assume that domain-specific operations analogous to our merge operation are available. Perhaps the closest effort to our own has been Pregoica *et al.*’s IceCube [30], and Edwards *et al.*’s Bayou [18], each of which exploit a range of specific algebraic properties of updates to distributed state. However, both systems must be explicitly adapted to specific application domains by the construction of domain-specific property oracles, or by mapping the application’s behavior down to a trivial update language. To the best of our knowledge, none of these areas have been explored in the context of a non-trivial state update language.

Update Sequencing. The use of distributed logs and publish/subscribe to apply a canonical order to updates has also been explored extensively by the distributed systems and database communities. Ellis *et al.* noted the relevance of sequencing to distributed concurrency control [19]. Eugster *et al.* identified the usefulness of sequencing updates to distributed collection types [20]. Domain specific applications of similar ideas can be found in work by Ostrowski and Birman [29], Weatherspoon *et al.* [33], and others.

Intent-Based Updates. The use of intent-based (*i.e.*, operational) updates appears frequently in database literature, especially in the context of distributed databases, where it is used to reduce communication overhead. Two concrete examples are Ceri and Widom’s Starburst [13], and Chang *et al.*’s BigTable [14].

6. REFERENCES

- [1] S. Abiteboul and N. Bidoit. Non first normal form relations: An algebra allowing data restructuring. *JCSS*, 33(3):361–393, 1986.
- [2] S. Abiteboul, D. Quass, J. McHugh, J. Widom, and J.L. Wiener. The Lorel query language for semistructured data. *JODL*, 1(1):68–88, 1997.
- [3] Serge Abiteboul and Catriel Beeri. The power of languages for the manipulation of complex values. *VLDBJ*, 4(4):727–794, October 1995.
- [4] Serge Abiteboul and Richard Hull. IFO: a formal semantic database model. *ACM TODS*, 12(4):525–565, November 1987.
- [5] Jiří Adámek, Mahdiah Haddadi, and Stefan Milius. From corecursive algebras to corecursive monads. In *CALCO*, pages 55–69, 2011.
- [6] B R Badrinath and Krithi Ramamritham. Performance evaluation of semantics-based multilevel concurrency control protocols. In *SIGMOD*, May 1990.
- [7] Adriana Balan and Alexander Kurz. On coalgebras over algebras. *Electron. Notes Theor. Comput. Sci.*, 264(2):47–62, August 2010.
- [8] M. Benedikt and J. Cheney. Semantics, types and effects for xml updates. *DBPL*, pages 1–17, 2009.
- [9] K. Beyer, V. Ercegovic, R. Gemulla, A. Balmin, M. Eltabakh, C.C. Kanne, F. Ozcan, and E.J. Shekita. Jaql: A scripting language for large scale semistructured data analysis. *PVLDB*, 4(12), 2011.
- [10] K. Beyer, V. Ercegovic, J. Rao, and E. Shekita. Jaql: A json query language. *URL: http://jaql.org*, 2009.
- [11] P. Buneman. Semistructured data. In *PODS*, pages 117–121, 1997.
- [12] P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theoretical Computer Science*, 149(1):3–48, 1995.
- [13] Stefano Ceri and Jennifer Widom. Production rules in parallel and distributed database environments. *PVLDB*, 1992.
- [14] F. Chang, J. Dean, S. Ghemawat, W.C. Hsieh, D.A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R.E. Gruber. Bigtable: A distributed storage system for structured data. *ACM TOCS*, 26(2):4, 2008.
- [15] J. Cheney. Lux: A lightweight, statically typed xml update language. *SIGPLAN*, 1060:25–36, 2007.
- [16] J. Cheney. Flux: functional updates for xml. *ACM SIGPLAN Notices*, 43(9):3–14, 2008.
- [17] S. Cluet, C. Delobel, C. Lécluse, and P. Richard. Reloop, an algebra based query language for an object-oriented database system. *Data Knowl. Eng.*, 5(4):333–352, October 1990.
- [18] W Keith Edwards, Elizabeth D Mynatt, and Karin Petersen. Designing and implementing asynchronous collaborative applications with Bayou. In *UIST*, 1997.
- [19] C A Ellis and S J Gibbs. Concurrency control in groupware systems. *SIGMOD*, 1989.
- [20] Patrick Th Eugster and Rachid Guerraoui. Distributed asynchronous collections: Abstractions for publish/subscribe interaction. *ECOOP*, 2000.
- [21] Ariel J Feldman, William P Zeller, Michael J Freedman, and Edward W Felten. SPORC: Group Collaboration using Untrusted Cloud Resources. In *OSDI*, 2010.
- [22] Martin Hyland and John Power. The category theoretic understanding of universal algebra: Lawvere theories and monads. *Electron. Notes Theor. Comput. Sci.*, 172:437–458, April 2007.
- [23] G. Jaeschke and H. J. Schek. Remarks on the algebra of non first normal form relations. In *PODS*, pages 124–138, 1982.
- [24] Christoph Koch. On the complexity of nonrecursive XQuery and functional query languages on complex values. *ACM TODS*, 31(4):1215–1256, December 2006.
- [25] K. Lellahi and V. Tannen. A calculus for collections and aggregates. In *Category Theory and Computer Science*, pages 261–280. Springer, 1997.
- [26] Zoran Majkic and Bhanu Prasad. Kleisli category and database mappings. *IJIDS*, 4(5):509–527, October 2010.
- [27] Ernie G. Manes. Implementing collection classes with monads. *Mathematical. Structures in Comp. Sci.*, 8(3):231–276, June 1998.
- [28] Gérald Oster, Pascal Urso, Pascal Molli, and Abdessamad Imine. Data Consistency for P2P Collaborative Editing. In *CSCW*, page 259, 2006.
- [29] Krzysztof Ostrowski and Ken Birman. Storing and accessing live mashup content in the cloud. *SIGOPS Review*, 44(2), April 2010.
- [30] Nuno Pregoica, Marc Shapiro, and Caroline Matheson. Semantics-based reconciliation for collaborative and mobile environments. *On The Move to Meaningful Internet . . .*, 2003.
- [31] M.A. Roth, H.F. Korth, and A. Silberschatz. Extended algebra and calculus for nested relational databases. *ACM TODS*, 13(4):389–417, 1988.
- [32] Marc Shapiro and Nuno Pregoica. Designing a commutative replicated data type. Technical report, CORR, October 2007.
- [33] Hakim Weatherspoon, Patrick Eaton, Byung-Gon Chun, and John Kubiatowicz. Antiquity: exploiting a secure log for wide-area distributed storage. In *EuroSys*, 2007.
- [34] William E Weihl. Commutativity-based concurrency control for abstract data types. *IEEE TC*, 37(12):1488–1505, 1988.